

Bayesian networks for supporting query processing over incomplete autonomous databases

Rohit Raghunathan · Sushovan De ·
Subbarao Kambhampati

Received: 28 August 2012 / Revised: 15 August 2013 / Accepted: 19 August 2013
© Springer Science+Business Media New York 2013

Abstract As the information available to naïve users through autonomous data sources continues to increase, mediators become important to ensure that the wealth of information available is tapped effectively. A key challenge that these information mediators need to handle is the varying levels of incompleteness in the underlying databases in terms of missing attribute values. Existing approaches such as QPIAD aim to mine and use Approximate Functional Dependencies (AFDs) to predict and retrieve relevant incomplete tuples. These approaches make independence assumptions about missing values—which critically hobbles their performance when there are tuples containing missing values for multiple correlated attributes. In this paper, we present a principled probabilistic alternative that views an incomplete tuple as defining a distribution over the complete tuples that it stands for. We learn this distribution in terms of Bayesian networks. Our approach involves mining/“learning” Bayesian networks from a sample of the database, and using it to do both imputation (predict a missing value) and query rewriting (retrieve relevant results with incompleteness on the query-constrained attributes, when the data sources are autonomous). We present empirical studies to demonstrate that (i) at higher levels of incompleteness, when multiple attribute values are missing, Bayesian networks do provide a significantly higher classification accuracy and (ii) the relevant possible answers retrieved by the queries reformulated using Bayesian networks

This research is supported by ONR grant N000140910032 and two Google research awards.

R. Raghunathan
Amazon, Seattle WA, USA
e-mail: rraghun@amazon.com

S. De (✉) · S. Kambhampati
Computer Science and Engineering, Arizona State University, Tempe AZ, USA
e-mail: sushovan@asu.edu

S. Kambhampati
e-mail: rao@asu.edu

provide higher precision and recall than AFDs while keeping query processing costs manageable.

Keywords Data cleaning · Bayesian networks · Query rewriting · Autonomous database

1 Introduction

As the popularity of the World Wide Web continues to increase, naïve users have access to more and more information in autonomous databases. Incompleteness in these autonomous sources is extremely commonplace. Such incompleteness mainly arises due to the way in which these databases are populated—by naïve users, or through (inaccurate) automatic extraction. Dealing with incompleteness in the databases requires tools for dealing with uncertainty. Previous attempts at dealing with this uncertainty by systems like QPIAD (Wolf et al. 2009) have mainly focused on using rule-based approaches, popularly known in the database community as Approximate Functional Dependencies (AFDs). The appeal of AFDs is due to the ease of specifying the dependencies, learning and reasoning with uncertainty. However, uncertain reasoning using AFDs adopts the certainty factors model, which assumes that the principles of locality and detachment (Russell and Norvig 2010) hold. But, these principles do not hold for uncertain reasoning and can lead to erroneous reasoning. As the levels of incompleteness in the information sources increases, the need for more scalable and accurate reasoning becomes paramount.

Full probabilistic reasoning avoids the traps of AFDs. Graphical models are an efficient way of doing full probabilistic reasoning. A Bayesian network is such a model, where direct dependencies between the variables in a problem are modeled as a directed acyclic graph, and the indirect dependencies can be inferred. As desired, Bayesian networks can model both causal and diagnostic dependencies. Using Bayesian networks for uncertain reasoning has largely replaced rule-based approaches in Artificial Intelligence. However, learning and inference on Bayesian networks can be computationally expensive which might inhibit their applications to handling incompleteness in autonomous data sources. In this paper, we consider if these costs can be handled without compromising on the improved accuracy offered by Bayesian networks, in the context of incompleteness in the autonomous databases.

1.1 Incompleteness in autonomous databases

Increasingly many of the autonomous web databases are being populated by automated techniques or by naïve users, with very little curation. For example, databases like autotrader.com are populated using automated extraction techniques by crawling the text classifieds and by car owners entering data through forms. Scientific databases such as CBioC (2013), also use similar techniques for populating the database. However, Gupta and Sarawagi (2006) have shown that these techniques are error prone and lead to incompleteness in the database in the sense that many of the attributes have missing values. Wolf et al. (2009) report that 99 % of the 35,000 tuples extracted from Cars Direct were incomplete. When the mediator has

privileges to modify the data sources, the missing values in these data sources can be completed using “imputation”, which attempts to fill in the missing values with the most likely value. As the levels of incompleteness in these data sources increase, it is not uncommon to come across tuples with multiple missing values. Effectively finding the most likely completions for these multiple missing values would require capturing the dependencies between them. A second challenge arises when the underlying data sources are autonomous, i.e., access to these databases are through forms, the mediator cannot complete the missing values with the most likely values. Therefore, mediators need to generate and issue a set of reformulated queries, in order to retrieve the relevant answers with missing values. Efficiency considerations dictate that the number of reformulations be kept low. In such scenarios, it becomes very important for mediators to send queries that not only retrieve results with a large fraction of relevant results (precision), but also a large number of relevant results (recall).

QPIAD & AFDs The QPIAD system addresses the challenges in retrieving relevant incomplete answers by learning the correlations between the attributes in the database as AFDs and the value distributions as Naïve Bayesian Classifiers. AFDs are rule-based methods for dealing with uncertainty. AFDs adopt the certainty factors model (Shortliffe 1976) which makes two strong assumptions:

1. **Principle of Locality:** Whenever there is a rule $A \rightarrow B$, given evidence of A , we can conclude B , regardless of the other rules and evidences.
2. **Principle of Detachment:** Whenever a proposition B is found to be true, the truth of B can be used regardless of how it was found to be true.

However, these two assumptions do not hold in the presence of uncertainty. When propagating beliefs, not only is it important to consider *all* the evidence but also their sources. Therefore, using AFDs for reasoning with uncertainty can lead to cyclic reasoning and fail to capture the correlations between multiple missing values. In addition to these shortcomings, the beliefs are represented using a Naive-Bayesian Classifier, which makes strong conditional independence assumptions, often leading to inaccurate values.

1.2 Overview of our approach

Given the advantages of Bayesian networks over AFDs, we investigate if replacing AFDs with Bayesian networks in QPIAD system, provides higher accuracy and while keeping the costs manageable. Learning and inference with Bayesian networks are computationally harder than AFDs. Therefore, the challenges involved in replacing AFDs with Bayesian networks include learning and using them to do both imputation and query rewriting by keeping costs manageable. We use BANJO software package (Hartemink et al. 2005) to learn the topology of the Bayesian network and use BNT (Murphy et al. 2001) and INFER.NET (Minka et al. 2010) software packages to do inference on them. Even though learning the topology for the Bayesian network from a sample of the database involves searching over the possible topologies, we found that high fidelity Bayesian networks could be learned from a small fraction of the database by keeping costs manageable (in terms of time spent in searching). Inference in Bayesian networks is intractable in

the worst case if the network is multiply connected, i.e., there is more than one undirected path between any two nodes in the network. We handle this challenge by using approximate inference techniques. Approximate inference techniques are able to retain the accuracy of exact inference techniques and keep the cost of inference manageable. We compare the cost and accuracy of using AFDs and Bayesian networks for imputing single and multiple missing values at different levels of incompleteness in test data.

We also develop new techniques for generating rewritten queries using Bayesian networks. The three challenges that are involved in generating rewritten queries are:

1. Selecting the attributes on which the new queries will be formulated. Selecting these attributes by searching over all the attributes becomes too expensive as the number of attributes in the database increases.
2. Determining the values to which the attributes in the rewritten query will be constrained to. The size of the domains of attributes in most autonomous databases is often large. Searching over each and every value can be expensive.
3. Most autonomous data sources have a limit on the number of queries to which it will answer. The rewritten queries that we generate should be able to carefully tradeoff precision with the throughput of the results returned.

We propose techniques to handle these challenges and compare them with AFD-based approaches in terms of precision and recall of the results returned.

Organization The rest of the paper is organized as follows—We begin with a discussion of related work, then in Section 3, we describe the problem setting and background. In Section 4, we discuss how Bayesian network models of autonomous databases can be learned by keeping costs manageable. In Section 5, we compare the prediction accuracy and cost of using Bayesian network and AFDs for imputing missing values. Next, in Section 6, we discuss how rewritten queries are generated using Bayesian networks and compare them with AFD-approaches for single and multi-attribute queries. Finally, we conclude in Section 7.

2 Related work

This work is a significant extension of the QPIAD system (Wolf et al. 2007, 2009), which also deals with incompleteness in databases. While the QPIAD system also learns attribute correlations, it does so using Approximate Functional Dependencies (AFDs) and uses Naive Bayesian Classifiers for representing value distributions and reformulating queries. Additionally, the QPIAD system can only handle missing values on a single attribute. In contrast, we use Bayesian network models learned from a sample of the database to represent attribute correlations and value distributions. We use the methods used in the QPIAD system as our baseline approach.

Completing missing values in databases using Bayesian networks has been addressed previously (Ramoni and Sebastiani 1997, 2001; Romero and Salmerón 2004; Fernández et al. 2012). Other methods have also been proposed to address learning from missing data, for example, Batista and Monard (2002) propose using k -Nearest Neighbor approach; Dempster et al. (1977) propose using an EM approach.

But most methods focus on completing the missing values so as to preserve the original data statistics so that other data mining techniques can be applied to it. We concentrate on retrieving relevant possible answers in the presence of missing values on single and multiple attributes. For example, Wu et. al. use association rules to impute the value of the missing attributes (Wu et al. 2004), whereas we use Bayesian networks to impute the value as well as retrieve results. In Section 5.2, we will show that using Bayesian networks is clearly superior to using rule-based imputation.

Like QPIAD, and other work on querying over incomplete databases, we too assume that the level of incompleteness in the database is small enough that it is possible to get a training sample that is mostly complete. Thus, we use techniques for learning from complete training data. If the training sample itself were to be incomplete, then we will need to employ expectation-maximization techniques during learning (Dempster et al. 1977).

Work on querying inconsistent databases usually focuses on fixing problems with the query itself (Muslea and Lee 2005; Nambiar and Kambhampati 2006). If the query has an empty result set, or if the query does not include all relevant keywords, it can be automatically augmented to fix those shortcomings. The objective of this work is to deal with shortcomings of the data—our query rewriting algorithms help retrieve useful tuples even in the presence of multiple missing values in them.

3 Problem setting & background

3.1 Overview of QPIAD

Since our main comparison is with the QPIAD system, we will provide a brief overview of its operation. Given a relation R , a subset X of its attributes, and a single attribute A of R , an approximate functional dependency (AFD) holds on a relation R , between X and A , denoted by, $X \rightsquigarrow A$, if the corresponding functional dependency $X \rightarrow A$ holds on all but a small fraction of tuples of R .

To illustrate how QPIAD works consider the query $Q : Body = SUV$ issued to Table 1. Traditional query processors will only retrieve tuples t_7 and t_9 . However, the entities represented by tuples t_8 and t_{10} are also likely to be relevant. The QPIAD system's aim is to retrieve tuples t_8 and t_{10} , in addition to t_7 and t_9 . In order to retrieve tuples t_8 and t_{10} it uses AFDs mined from a sample of the database. For example, an AFD $Model \rightsquigarrow Body$ may be mined for the fragment of the cars database shown in Table 1. This indicates that the value of a car's *Model* attribute often (but not always) determines the value of its *Body* attribute. These rules are used to retrieve relevant incomplete answers.

When the mediators have access privileges to modify the database, AFDs are used along with Naive Bayesian Classifiers to fill in the missing values as a simple classification task and then traditional query processing will suffice to retrieve relevant answers with missing values. However, in more realistic scenarios, when such privileges are not provided, mediators generate a set of rewritten queries and send to the database, in addition to the original user query. According to the AFD mentioned above and tuple t_7 retrieved by traditional query processors, a rewritten query $Q'_1 : \sigma_{Model=Santa}$ may be generated to retrieve t_8 . Similarly $Q'_2 : \sigma_{Model=MDX}$ may be generated which will retrieve t_{10} .

Table 1 A fragment of a car database

ID	Make	Model	Year	Body	Mileage
1	Audi	Null	Null	Sedan	20000
2	Audi	A8	Null	Sedan	15000
3	BMW	745	2002	Sedan	40000
4	Audi	Null	2005	Sedan	20000
5	Audi	A8	2005	Sedan	20000
6	BMW	645	1999	Convrt	Null
7	Hyundai	Santa	1990	SUV	45000
8	Hyundai	Santa	1993	Null	40000
9	Acura	MDX	1990	SUV	30000
10	Acura	MDX	1990	Null	12000

Multiple rules can be mined for each attribute, for example, the mileage and year of a car might determine the body style of the car. So a rule $\{Year, Mileage\} \rightsquigarrow \{Body\}$ could be mined. Each rule has a confidence associated with it, which specifies how accurate the determining set of an attribute's AFD is in predicting it. The current QPIAD system uses only the highest confidence AFD¹ of each attribute for imputation and query rewriting. In addition, it only aims to retrieve relevant incomplete answers with atmost one missing value on query-constrained attributes.

To illustrate the shortcomings of AFDs, consider a query $Q : \sigma_{Model=A8 \wedge Year=2005}$ issued to the fragment of the car database shown in Table 1. When the mediator has modification privileges, the missing values for attributes *Model* and *Year* can be completed with the most likely values, before returning the answer set. Using AFDs to predict the missing values in tuple t_1 , ignores the correlation between the *Model* and *Year*; predicting them independently. Substituting the value for missing attribute *Year* in tuple t_2 using just the highest confidence rule as is done in QPIAD, often leads to inaccurate propagation of beliefs as the other rules are ignored. When the mediator does not have privileges to modify the database, a set of rewritten queries are generated and issued to the database to retrieve the relevant uncertain answers. Issuing Q to the database fragment in Table 1 retrieves t_5 . The rewritten queries generated by methods discussed in QPIAD retrieve tuples t_2 and t_4 . However, it does not retrieve tuple t_1 , but it is highly possible that the entity represented by it is relevant to the user's query.

3.2 Bayesian networks

A Bayesian network (Pearl 1988) is a graphical representation of the probabilistic dependencies between the variables in a domain. The generative model of a relational database can be represented using a Bayesian network, where each node in the network represents an attribute in the database. The edges between the nodes represent direct probabilistic dependencies between the attributes. The strength of these probabilistic dependencies are modeled by associating a conditional probability distribution (CPD) with each node, which represents the conditional probability of a variable, given the combination of values of its immediate parents. A Bayesian

¹The actual implementation of QPIAD uses a variant to the highest confidence AFD for some of the attributes. For details we refer the reader to Wolf et al. (2009).

network is a compact representation of the full joint probability distribution of the nodes in the network. The full joint distribution can be constructed from the CPDs in the Bayesian network. Given the full joint distribution, any probabilistic query can be answered. In particular, the probability of any set of hypotheses can be computed, given any set of observations, by conditioning and marginalizing over the joint distribution. Since the semantics of Bayesian networks are in terms of the full joint probability distribution, inference using them considers the influence of all variables in the network. Therefore, Bayesian networks, unlike AFDs, do not make the Locality and Detachment assumptions.

The Markov blanket of a node X in the Bayesian network is the set of nodes comprising its parents, its children, and its children's other parents. If the value of all the nodes in the Markov blanket of X are given as evidence, then the value of X is independent of all other nodes in the network (Pearl 1988). The junction-tree algorithm (Jensen et al. 2006), also known as a clique tree, is a message passing algorithm that efficiently evaluates the posterior distribution on a Bayesian network. Gibb's sampling (Bishop et al. 2006) is a method of approximate inference on a Bayesian network that relies on Markov sampling of the network.

4 Learning Bayesian network models

In this section we discuss how we learn the topology and parameters of the Bayesian network by keeping costs manageable. We learn the generative model of two databases—The first dataset is extracted from *Cars.com* (Cars.com 2013) with the schema *Cars*(*Model*, *Year*, *Body*, *Make*, *Price*, *Mileage*). We call it the 'car database' in this paper. The second database is the adult database consisting of 15000 tuples obtained from UCI data repository (Frank and Asuncion 2010) with the schema *Adult*(*WorkClass*, *Occupation*, *Education*, *Sex*, *HoursPerWeek*, *Race*, *Relationship*, *NativeCountry*, *MaritalStatus*, *Age*). The car database has 55,000 tuples. We extract a fragment of 8,000 tuples from this dataset to learn the generative model. The adult database has 15,000 tuples, and we use the entire dataset for learning. Tables 2 and 3 describe the schema and the domain sizes of the attributes in the two databases. The first row in these table corresponds to the scenario where a mediator system is accessing the database. In such a scenario, the system does not have complete knowledge about the domain, and is therefore able to access only a fraction of the data. The attributes with continuous values are discretized and used as categorical attributes. *Price* and *Mileage* attributes in the cars database are discretized by rounding off to the nearest five thousand. In the adult database attributes *Age* and *Hours Per Week* are discretized to the nearest multiple of five.

The structure of the Bayesian network is learned from a complete sample of the autonomous database. We use the BANJO package (Hartemink et al. 2005) as a black box for learning the structure of the Bayesian network. To keep the learning

Table 2 Domain size of attributes in car database

Database	Year	Model	Make	Price	Mileage	Body
Cars-8000-20(Mediator)	9	38	6	19	17	5
Cars-8000-100(Complete)	12	41	6	30	20	7


Table 3 Domain size of attributes in adult database

Database	Age	Work class	Education	Marital status	Occupation	Relationship	Race	Sex	Hours per week	Native country
Adult-15000-20(Mediator)	8	7	16	7	14	6	5	2	10	37
Adult-15000-100(Complete)	8	7	16	7	14	6	5	2	10	40

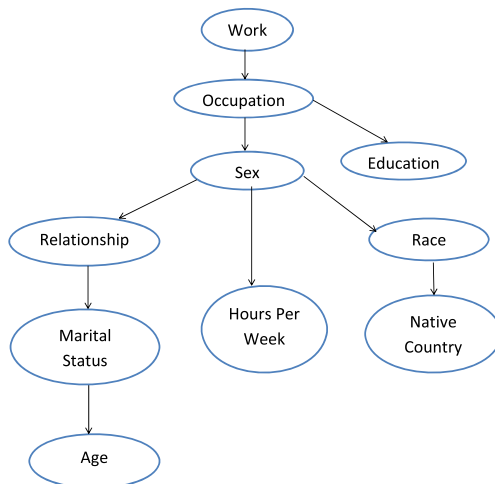
costs manageable we constrain nodes to have at most two parents. In cases where there are more than two attributes directly correlated to an attribute, these attributes can be modeled as children. There is no limit on the number of children a node can have. Figure 1a shows the structure of a Bayesian network learned for the Cars database and Fig. 1b for the adult database. We used samples of sizes varying from 5–20 % of the database and found that the structure of the highest scoring network remained the same.

The BANJO settings used were: searcher choice = Simulated Annealing, proposer choice = Random Local Move (Addition, deletion, or reversal of an edge in the

Fig. 1 Bayesian networks learned from a sample of the data. The size of the sample was varied from 5 to 20 %; the same structure was observed in all cases



(a) *Cars.com* dataset



(a) adult database

current network, selected at random), evaluator choice = BDe (Heckerman et al. 1995), decider choice = Metropolis (A Metropolis-Hastings stochastic decision mechanism, where any network with a higher score is accepted, and any with a lower score is accepted with a probability based on a system parameter known as the “temperature”).

We also experimented with different time limits for the search, ranging from 5–30 min. We did not see any change in the structure of the highest confidence network.

Bayesian network inference is used in both imputation and query rewriting tasks. Imputation involves substituting the missing values with the most likely values, which involves inference. Exact inference in Bayesian networks is NP-hard (Cooper 1990) in the worst case, if the network is multiply connected. Therefore, to keep query processing costs manageable we use approximate inference techniques. In our experiments described in Section 5, we found that using approximate inference techniques retains the accuracy edge of exact inference techniques, while keeping the prediction costs manageable. We use the BNT package (Murphy et al. 2001) for doing inference on the Bayesian Network² for the imputation task. We experimented with various exact inference engines that BNT offers and found the junction-tree engine to be the fastest. While querying multiple variables, junction tree inference engine can be used only when all the variables being queried form a clique. When they do not form a clique we use the variable elimination algorithm (Jensen and Nielsen 2007).

5 Imputation using Bayesian networks

In this section we compare the prediction accuracy and cost of Bayesian networks versus AFDs for imputing single and multiple missing values when there is incompleteness in test data. The prediction accuracy will be measured in controlled experiments where the ground truth is known, and the cost is determined as the time taken by the algorithm to generate the result. When the mediator has privileges to modify the underlying autonomous database, the missing values can be substituted with the most probable value. Imputation using Bayesian networks first computes the posterior of the attribute that is to be predicted given the values present in the tuple and completes the missing value with the most likely value given the evidence. When predicting multiple missing values, the joint posterior distribution over the missing attributes are computed and the values with the highest probability are used for substituting the missing values. Computing the joint probability over multiple missing values captures the correlations between the missing values, which gives Bayesian networks a clear edge over AFDs. In contrast, imputation using AFDs uses the AFD with the highest confidence for each attribute for prediction. If an attribute in the determining set of an AFD is missing, then that attribute is first predicted using other AFDs (chaining), before the original attribute can be predicted. The most likely value for each attribute is used for completing the missing value. When multiple missing values need to be predicted, each value is predicted independently.

²In this prototype, we manually transferred the output of the BANJO module to the BNT module. In future systems, we will integrate them programmatically.

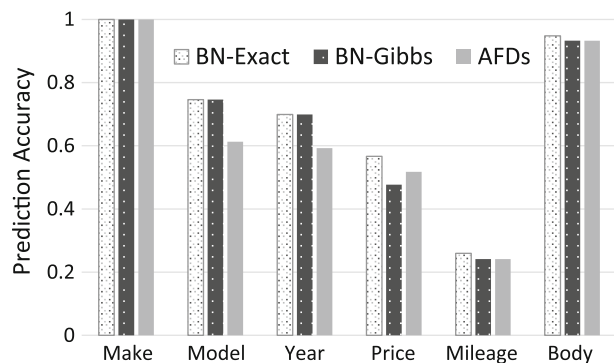
Before discussing the experimental results, we note that an implicit assumption for our method to work is that the incompleteness in the data is random. In more precise statistical terms, we assume that our data attributes are *missing at random* (Heitjan and Basu 1996). While this assumption does hold by-and-large, it is certainly not guaranteed to always hold in real world data. In used car data, for example, the owners of cars with known defects (e.g. Yugo), may well opt to not specify the make. If this happens systematically and completely, then our system will not be able to learn how make being Yugo will be correlated with any other attribute/value combination.

We use the car and adult databases described in the previous section. We compare AFD approach used in QPIAD which uses Naïve Bayesian Classifiers to represent value distributions with exact and approximate inference in Bayesian networks. We call exact inference in Bayesian network as BN-Exact. We use Gibbs sampling as the approximate inference technique, which we call BN-Gibbs. For BN-Gibbs, the probabilities are computed using 250 samples. For the imputation experiments, the data was divided into a training set and a test set. In the test set, randomly chosen values from the attribute being tested were removed.

5.1 Imputing single missing values

Our experiments show that prediction accuracy using Bayesian networks is higher than AFDs for attributes which have multiple high confidence rules. Approaches for combining multiple rules for classification have been shown to be ineffective by Khatri (2006). Since there is no straightforward way for propagating beliefs using multiple AFDs, only the AFD with the highest confidence is used for propagating beliefs. This method, however, fails to take advantage of additional information that the other rules provide. Bayesian networks, on the other hand, systematically combine evidences from multiple sources. Figure 2 shows the prediction accuracy in the presence of a single missing value for each attribute in the Cars database. We notice that there is a statistically significant difference in prediction accuracies by the two approaches for the attributes *Model* and *Year*. There are multiple rules that are mined for these two attributes but using just the rule with highest confidence, ignores the influence of the other available evidence, which affects the prediction accuracy.

Fig. 2 Single attribute prediction accuracy (cars)



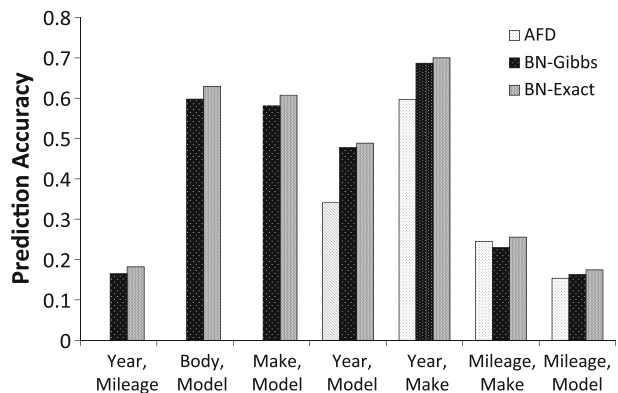
5.2 Imputing multiple missing values

In most real-world scenarios, however, the number of missing values per tuple is likely to be more than one. The advantage of using a more general model like Bayesian networks becomes even more apparent in these cases. Firstly, AFDs cannot be used to impute all combinations of missing values, this is because when the determining set of an AFD contains a missing attribute, then the value needs to be first predicted using a different AFD by chaining. While chaining, if we come across an AFD containing the original attribute to be predicted in its determining set, then predicting the missing value becomes impossible. When the missing values are highly correlated, AFDs often get into such cyclic dependencies. In Fig. 3 we can see that the attribute pairs *Year-Mileage*, *Body-Model* and *Make-Model* cannot be predicted by AFDs. As the number of missing values increases, the number of combinations of missing values that can be predicted reduces. In our experiments with the Cars database, when predicting three missing values, only 9 out of the 20 possible combinations of missing values could be predicted.

On the other hand, Bayesian networks can predict the missing values regardless of the number and combination of values missing in a tuple. Secondly, while predicting the missing values, Bayesian networks compute the joint probability distribution over the missing attributes which allows them to capture the correlations between the attributes. In contrast, prediction using AFDs, which use a Naïve Bayesian Classifier to represent the value distributions, predict each of the missing attributes independently, ignoring the interactions between them. The attribute pair *Year-Model* in Fig. 3 shows that the prediction accuracy is significantly higher when correlations between the missing attributes are captured. We also observe that in some cases, when the missing values are D-separated (Geiger et al. 1990) given the values for other attributes in the tuple, the performance of AFDs and Bayesian networks is comparable. In Fig. 3, we can see that the prediction accuracy for *Mileage-Make* and *Mileage-Model* are comparable for all the techniques since attributes are D-separated given the other evidence in the tuple. However, the number of attributes that are D-separated is likely to decrease with increase in incompleteness in the databases.

There is an interesting trade-off between the time taken and the accuracy of imputation when dealing with tuples with 2 or more missing attributes. Suppose

Fig. 3 Multiple attribute prediction accuracy (cars)



that the attributes that are missing are $X = \{x_1, x_2, \dots, x_n\}$ and the attributes that are known are $A = \{a_1, a_2, \dots, a_m\}$. Then one approach would be to find those attribute values for x_i that maximize $P(x_i|A)$ individually. This would be the “most likely estimation” approach, and the fastest method; however, this is also the least accurate since it ignores all interactions among the attributes in X . A slightly more accurate approach would be a “greedy search” approach, where the value for one of the variables in X is (i) found using $\arg \max(x_i|A)$, and then (ii) set as evidence. This process is repeated for the remaining unknown attributes. The imputation of those attributes is, therefore, informed by $A \cup x_i$. An even more accurate approach is to use “most probable explanation”, where we find the entire set of the attributes X that maximizes the joint conditional probability $P(X|A)$. This is the approach that we use in this paper.

5.3 Prediction accuracy with increase in incompleteness in test data

As the incompleteness in the database increases, not only does the number of values that need to be predicted increase, but also the evidence for predicting these missing values reduces. We compared the performance of Bayesian networks and AFDs as the incompleteness in the autonomous databases increases. We see that the prediction accuracy of AFDs drops faster with increase in incompleteness. This is because the chaining required for predicting missing values using AFDs increases which in turn increases the chances of getting into cyclic dependencies. Also, when an attribute has multiple AFDs, propagating beliefs using just one rule and ignoring the others, often violates the principles of detachment and locality (Heckerman 1992) impacting the prediction accuracy.

On the other hand, Bayesian networks, being a generative model, can systematically infer the values of any set of attributes given the evidence of any other set. Therefore, as the incompleteness of the database increases, the prediction accuracy of Bayesian Networks will be significantly higher than that of AFDs. Figure 4 shows the prediction accuracy of AFDs and Bayesian networks when single and multiple attributes need to be predicted in the car and adult databases. We see that both Bayesian networks approaches have a higher prediction accuracy than AFDs at all levels of incompleteness.

5.4 Time taken for imputation

We now compare the time taken for imputing the missing values using AFDs, exact inference (junction tree) and Gibbs sampling (250 samples) as the number of missing values in the autonomous database increases. Table 4 reports the time taken to impute a Cars database with 5,479 tuples and Fig. 4 shows the accuracy. In Fig. 4 we can clearly see that the AFD based method is outperformed by both the Gibbs sampling based inference method as well as the Exact Inference method. Graphs (a–b) use the used car database and graphs (c–d) use the census dataset. We can see from the curves that using approximate inference is not much less accurate than using exact inference both in the case of single attribute imputation (a, c) and two attribute

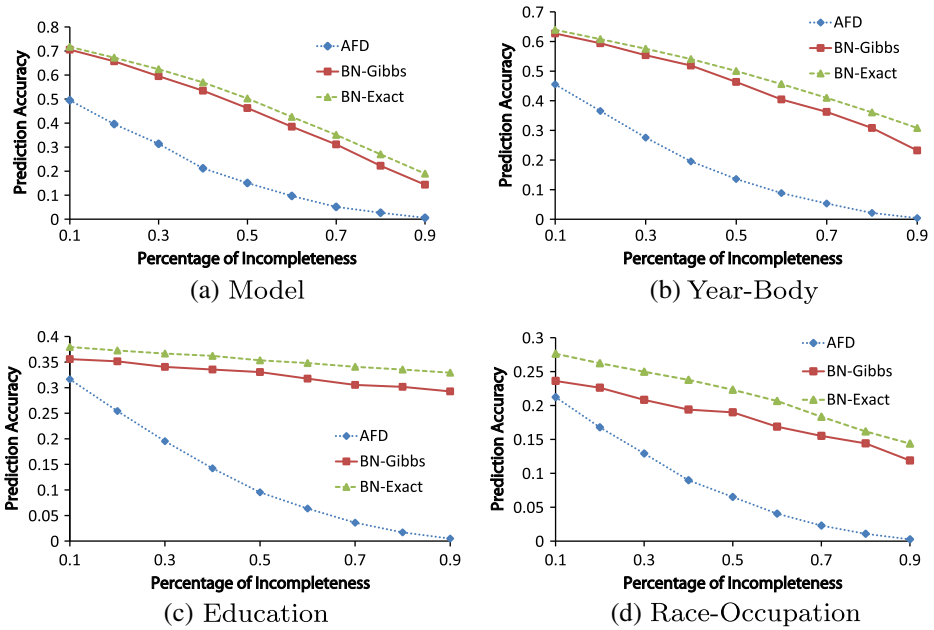


Fig. 4 Prediction Accuracy with increase in percentage of incompleteness in test data for various missing attributes. **a–c** Car database, **d** adult database

imputation (b, d). The difference is about 4 percentage points in the worst case, so we can say that the preferred method for most applications is Gibbs sampling.

Table 4 shows the time taken by each of these methods as the percentage of incompleteness varies. Note that the time taken by the AFD method reduces with higher incompleteness because there are fewer AFDs learned from the data.

Table 4 Time taken for predicting 5,479 tuples by AFDs, BN-Gibbs (250 samples) and BN-exact in seconds

Percentage of incompleteness (%)	Time taken for AFD (s)	Time taken by BN-Gibbs (s)	Time taken by BN-exact (s)
0	0.271	44.46	16.23
10	0.267	47.15	44.88
20	0.205	52.02	82.52
30	0.232	54.86	128.26
40	0.231	56.19	182.33
50	0.234	58.12	248.75
60	0.232	60.09	323.78
70	0.235	61.52	402.13
80	0.262	63.69	490.31
90	0.219	66.19	609.65

6 Query rewriting with Bayesian networks

In information integration scenarios when the underlying data sources are autonomous, missing values cannot be completed using imputation. Our goal is to retrieve *all* relevant answers to the user's query, including tuples which are relevant, but have missing values on the attributes constrained in the user's query. Since query processors are allowed read-only access to these databases, the only way to retrieve the relevant answers with missing values on query-constrained attributes is by generating and sending a set of reformulated queries that constrain other relevant attributes. We describe two techniques—BN-All-MB and BN-Beam—for retrieving such relevant incomplete results using Bayesian networks.

6.1 Generating rewritten queries

Table 5 shows a different fragment of the database shown in Table 1. We will use this fragment to explain our approach. Notice that tuples t_2, t_3 have one missing (null) value and tuples t_5, t_6, t_7, t_8 have two missing values. To illustrate query rewriting when a single attribute is constrained in the query, consider a query(Q) $\sigma_{\text{Body}=\text{Sedan}}$. First, the query Q is issued to the autonomous database and all the certain answers which correspond to tuples t_1, t_3, t_4 and t_5 in the Table 5 are retrieved. This set of certain answers forms the *base result set*. However, tuple t_2 , which has a missing value for *Body* (possibly due to incomplete extraction or entry error), is likely to be relevant since the value for *Body* should have been *Sedan* had it not been missing. In order to determine the attributes and their values on which the rewritten queries need to be generated, we use the Bayesian network learned from the sample of the autonomous database.

Using the same example, we now illustrate how rewritten queries are generated. First, the set of certain answers which form the base result set are retrieved and returned to the user. The attributes on which the new queries are reformulated consist of all attributes in the *Markov blanket* of the original query-constrained attribute. We consider all attributes in the Markov blanket while reformulating queries because given the values of these attributes, the original query-constrained attribute is dependent on no other attribute in the Bayesian network. From the learned Bayesian network shown in Fig. 1a, the Markov blanket of the attribute *Body* consists of the attributes $\{\text{Year}, \text{Model}\}$. The value that each of the attributes in the rewritten query can be constrained to is limited to the distinct value combinations

Table 5 A Fragment of a car database

ID	Make	Model	Year	Body	Mileage
1	Audi	A8	2005	Sedan	20000
2	Audi	A8	2005	Null	15000
3	Acura	tl	2003	Sedan	Null
4	BMW	745	2002	Sedan	40000
5	Null	745	2002	Sedan	Null
6	Null	645	1999	Convrt	Null
7	Null	645	1999	Coupe	Null
8	Null	645	1999	Convrt	Null
9	BMW	645	1999	Coupe	40000
10	BMW	645	1999	Convrt	40000

Algorithm 1 Algorithm for BN-All-MB

Let $R(A_1, A_2, \dots, A_n)$ be a database relation. Suppose $MB(A_p)$ is the set of attributes in the Markov blanket of attribute A_p . A query $Q: \sigma_{A_p=v_p}$ is processed as follows

1. Send Q to the database to retrieve the base result set $RS(Q)$. Show $RS(Q)$, the set of certain answers, to the user.
 2. Generate a set of new queries Q' , order them, and send the most relevant ones to the database to retrieve the extended result set $\widehat{RS}(Q)$ as relevant possible answers of Q . This step contains the following tasks.
 - (a) *Generate Rewritten Queries.* Let $\pi_{(MB(A_p))}(RS(Q))$ be the projection of $RS(Q)$ onto $MB(A_p)$. For each distinct tuple t_i in $\pi_{(MB(A_p))}(RS(Q))$, create a selection query Q'_i in the following way. For each attribute A_x in $MB(A_p)$, create a selection predicate $A_x=t_i.v_x$. The selection predicates of Q'_i consist of the conjunction of all these predicates
 - (b) *Select the Rewritten Queries.* For each rewritten query Q'_i , compute the estimated precision and estimated recall using the Bayesian network as explained earlier. Then order all Q'_i s in order of their F-Measure scores and choose the top- K to issue to the database.
 - (c) *Order the Rewritten Queries.* The top- K Q'_i s are issued to the database in the decreasing order of expected precision.
 - (d) *Retrieve extended result set.* Given the ordered top- K queries $\{Q'_1, Q'_2, \dots, Q'_K\}$ issue them to the database and retrieve their result sets. The union of result sets $RS(Q'_1), RS(Q'_2), \dots, RS(Q'_K)$ is the extended result set $\widehat{RS}(Q)$.
-

for each attribute in the base result set. This is because, the values that the other attributes take are highly likely to be present in relevant incomplete tuples. This tremendously reduces the search effort required in generating rewritten queries, without affecting the recall too much. However, at higher levels of incompleteness, this might have a notable impact on recall, in which case, we would search over the values in the entire domain of each attribute. Continuing our example, when the query Q is sent to the database fragment shown in Table 5, tuples t_1, t_3, t_4 and t_5 are retrieved. The values over which the search is performed for *Model* is $\{A8, tl, 745\}$ and for *Year* is $\{2002, 2003, 2005\}$.

Some of the rewritten queries that can be generated by this process are

$$Q'_1: \sigma_{Model=A8 \wedge Year=2005}, Q'_2: \sigma_{Model=tl \wedge Year=2003} \text{ and}$$

$$Q'_3: \sigma_{Model=745 \wedge Year=2002}.$$

Each of these queries differ in the number of results that they retrieve and the fraction of retrieved results that are relevant. An important issue here is to decide which of these queries should be issued to the autonomous database and in which order. If we are allowed to send as many queries as we want, ordering the queries in terms of their *expected precision* would obviate the need for ranking the relevant possible results once they are retrieved. This is because the probability that the missing value in a tuple is exactly the value the user is looking for is the same as the expected precision of the query that retrieves the tuple. However, limits are often

imposed on the number of queries that can be issued to the autonomous database. These limits could be due to network or processing resources of the autonomous data sources. Given such limits, the precision of the answers need to be carefully traded off with selectivity (the number of results returned) of the queries. One way to address this challenge is to pick the top- K queries based on the *F-measure* metric (Manning et al. 2008), as pointed out by Wolf et al. F-measure is defined as the weighted harmonic mean of precision (P) and recall (R) measures :

$$\frac{(1 + \alpha) \cdot P \cdot R}{\alpha \cdot P + R}$$

For each rewritten query, the F-measure metric is evaluated in terms of its *expected precision* and *expected recall*. The latter which can be computed from the *expected selectivity*. Expected precision is computed from the Bayesian network and expected selectivity is computed the same way as computed by the QPIAD system, by issuing the query to the sample of the autonomous database. For our example, the expected precision of the rewritten query $\sigma_{Model=A8 \wedge Year=2005}$ can be computed as the $P(\text{Body}=\text{Sedan} \mid \text{Model}=A8 \wedge \text{Year}=2005)$ which is evaluated by inference on the learned Bayesian network. Expected selectivity is computed as $\text{SmplSel}(Q) \cdot \text{SmplRatio}(R)$, where $\text{SmplSel}(Q)$ is the sample selectivity of the query Q , which is the fraction of the tuples in the sample returned by the query and $\text{SmplRatio}(R)$ is the ratio of the original database size to the size of the sample. For example, if the original database has 1 million tuples, and our sample database has 10,000 tuples, then the $\text{SmplRatio}(R)$ is $10^4/10^6 = 0.01$. If the query results have 1,000 tuples in them, then $\text{SmplSel}(Q)$ is $10^3/10^4 = 0.1$. We send queries to the original database and its sample offline and use the cardinalities of the result sets to estimate the ratio.

We refer to this technique for generating rewritten queries by constraining all attributes in the Markov blanket as BN-All-MB. In Section 6.2.1 we compare the performance of BN-All-MB and AFD approaches in retrieving uncertain relevant tuples. However, the issue with constraining all attributes in the Markov blanket is that its size could be arbitrarily large. Since the Markov blanket comprises the children, the parents, and the children's other parent nodes, the number of attributes constrained in the rewritten queries becomes very large. This will reduce the throughput of the queries significantly. As we mentioned earlier, in cases where the autonomous database has a limit on the number of queries to which it will respond, we need to carefully trade off precision of the rewritten queries with their throughput. BN-All-MB and AFD approaches decide upfront the attributes to be constrained and search only over the values to which the attributes will be constrained. Both these techniques try to address this issue by using the F-measure metric to pick the top- K queries for issuing to the database—all of which have the same number of attributes constrained. A more effective way to trade off precision with the throughput of the rewritten queries is by making an “online” decision on the number of attributes to be constrained. We propose a technique, BN-Beam which searches over the Markov blanket of the original query-constrained attribute, and picks the best subset of the attributes with high precision and throughput.

6.1.1 Generating rewritten queries using BN-beam

We now describe BN-Beam, our technique for generating rewritten queries which finds a subset of the attributes in the Markov blanket of the query-constrained attribute with high precision and throughput. To illustrate how rewritten queries are generated using BN-Beam, consider the same query(Q) $\sigma_{Body=Sedan}$. First, Q is sent to the database to retrieve the base result set. We consider the attributes in the Markov blanket of the query-constrained attribute to be the potential attributes on which the new queries will be formulated. We call this set the *candidate attribute set*.

For query Q , the candidate attribute set consists of attributes in the Markov blanket of attribute *Body* which consists of attributes $\{Model, Year\}$ for the Bayesian Network in Fig. 1a. Once the candidate attribute set is determined, a beam search with a beam width, K and level, L , is performed over the distinct value combinations in the base result set of the attributes in the candidate attribute set. For example, when the query Q is sent to the database fragment shown in Table 5, tuples t_1, t_3, t_4 and t_5 are retrieved. The values over which the search is performed for *Model* is $\{A8, tl, 745\}$ and for *Year* is $\{2002, 2003, 2005\}$. Starting from an empty rewritten query, the beam search is performed over multiple levels, looking to expand the partial query at the previous level by adding an attribute-value to it. For example, at the first level of the search five partial rewritten queries: $\sigma_{Model=745}$, $\sigma_{Model=A8}$, $\sigma_{Model=tl}$, $\sigma_{Year=2002}$ and $\sigma_{Year=2003}$ may be generated. An important issue here is to decide which of the queries should be carried over to the next level of search. Since there is a limit on the number of queries that can be issued to the autonomous database and we want to generate rewritten queries with high precision and throughput while keeping query processing costs low, we pick the top- K queries based on the F-measure metric, as described earlier. The advantage of performing a search over both attributes and values for generating rewritten queries is that there is much more control over the throughput of the rewritten queries as we can decide how many attributes will be constrained.

The top- K queries at each level are carried over to the next level for further expansion. For example, consider query $\sigma_{Model=745}$ which was generated at level one. At level two, we try to create a conjunctive query of size two by constraining the other attributes in the candidate attribute set. Say we try to add attribute *Year*, we search over the distinct values of *Year* in the base set with attribute *model* taking the value 745. At each level i , we will have the top- K queries with highest F-measure values with i or fewer attributes constrained. The top- K queries generated at the Level L are sorted based on expected precision and sent to the autonomous database in that order to retrieve the relevant possible answers. We now describe the BN-Beam algorithm for generating rewritten queries for single-attribute queries.

In step 2(d), it is important to remove duplicates from $\widehat{RS}(Q)$. Since rewritten queries may constrain different attributes, the same tuple might be retrieved by different rewritten queries. For example, consider two rewritten queries- $Q'_1: \sigma_{Model=A8}$ and $Q'_2: \sigma_{Year=2005}$, that can be generated at level one for the same user query Q , that aims to retrieve all *Sedan* cars. All *A8* cars manufactured in 2005 will be returned in the answer sets of both queries. Therefore, we need to remove all duplicate tuples.

Algorithm 2 Algorithm for BN-Beam

Let $R(A_1, A_2, \dots, A_n)$ be a database relation. Suppose $MB(A_p)$ is the set of attributes in the Markov blanket of attribute A_p . All the steps in processing a query $Q: \sigma_{A_p=v_p}$ is the same as described for BN-All-MB except step 2(a) and 2(d).

- 2(a) Generate Rewritten Queries. A beam search is performed over the attributes in $MB(A_p)$ and the value for each attribute is limited to the distinct values for each attribute in $RS(Q)$. Starting from an empty rewritten query, a partial rewritten query (PRQ) is expanded, at each level, to add an attribute-value pair from the set of attributes present in $MB(A_p)$ but not added to the partial rewritten query already. The queries with top- K values for F-measure scores, computed from the estimated precision and estimated recall computed from the sample, are carried over to the next level of the search. The search is repeated over L levels.
- 2(d) Post-filtering. Remove the duplicates in $\widehat{RS}(Q)$.

6.1.2 Handling multi-attribute queries

Retrieving relevant uncertain answers for multi-attribute queries has been only superficially addressed in the QPIAD system. It attempts to retrieve only uncertain answers with missing values on any one of the multiple query-constrained attributes. Here we describe how BN-All-MB and BN-Beam can be extended to retrieve tuples with missing values on multiple query-constrained attributes.

BN-All-MB: The method described to handle single-attribute queries using BN-All-MB can be easily generalized to handle multi-attribute queries. The rewritten queries generated will constrain every attribute in the union of the Markov blanket of the constrained attributes.

BN-Beam: Similarly, using BN-Beam to handle multi-attribute queries is simple extension of the method described for single-attribute queries. The candidate attribute set consists of the union of the attributes in the Markov blanket of each query-constrained attribute.

To illustrate how new queries are reformulated to retrieve possibly relevant answers with multiple missing values on query-constrained attributes, consider an example query $\sigma_{Make=BMW \wedge Mileage=40000}$ sent to database fragment in Table 5. First, this query retrieves the base result set which consists of tuples t_4, t_9, t_{10} . The set of candidate attributes on which the new queries will be formulated is obtained by the union of attributes in the Markov blanket of the query-constrained attributes. For the learned Bayesian network shown in Fig. 1a, this set consists of $\{Model, Year\}$. Once the candidate attribute set is determined, a beam search with a beam width, K , is performed similar to the case when a single attribute is constrained. At the first level of the search some of the partial rewritten queries that can be generated are $\sigma_{Model=745}$, $\sigma_{Model=645}$ and $\sigma_{Year=1999}$. The top- K queries with highest F-measure values are carried over to the next level of the search. The top- K queries generated at the Level L are sent to the autonomous database to retrieve the relevant possible answers.

6.2 Empirical evaluation of query rewriting

The aim of the experiments reported in this section is to compare the precision and recall of the relevant uncertain results returned by rewritten queries generated by AFDs and Bayesian networks for single and multi-attribute queries. We use the datasets as described in Section 4 and partition them into test and training sets. 15 % of the tuples are used as the training set. The training set is used for learning the topology and parameters of the Bayesian network and AFDs. It is also used for estimating the expected selectivity of the rewritten queries. We use the Expectation Propagation inference algorithm (Minka 2001) (with 10 samples) available in Infer.NET software package (Minka et al. 2010) for carrying inference on the Bayesian network.

In order to evaluate the relevance of the answers returned, we create a copy of the test dataset which serves as the *ground truth* dataset. We further partition the test data into two halves. One half is used for returning the certain answers, and in the other half all the values for the constrained attribute(s) are set to null. Note that this is an aggressive setup for evaluating our system. This is because typical databases may have less than 50 % incompleteness and even the incompleteness may not be on the query-constrained attribute(s). The tuples retrieved by the rewritten queries from the test dataset are compared with the ground truth dataset to compute precision and recall. Since the answers returned by the certain result set will be the same for all techniques, we consider only uncertain answers while computing precision and recall.

6.2.1 Comparison of rewritten queries generated by AFDs and BN-All-MB

Figure 5 shows the precision-recall curve for queries on attribute *Make* in the car database. The size of the Markov blanket and the determining set is one for attribute *Make*. We note that there is no difference in the quality of the results returned by AFDs and BN-All-MB in this case (see Fig. 5). Next, we compare the quality of the results returned by AFDs and BN-All-MB when the size of the Markov blanket and determining set of the AFD of the constrained attribute is greater than one. Figure 5 shows the precision and recall curves for the queries issued to car and adult databases. For the query on the adult database, we found that the order in which the rewritten queries were ranked were exactly the same. Therefore, we find that the precision-recall curves of both the approaches lie one on top of the other. For the queries issued to the car database, we find that there are differences in the order in which the rewritten queries are issued to the database. However, we note that there is no clear winner. The curves lie very close to each other, alternating as the number of results returned increases. Therefore the performance of AFDs and BN-All-MB is comparable for single-attribute queries.

6.2.2 Comparison of rewritten queries generated by BN-All-MB and BN-beam

Figure 6a shows the increase in recall of the results for three different values of α in the F-measure metric when ten queries can be issued to the database. We refer to results for different values for α for BN-Beam as BN-Beam- α (substitute α with its value) and show a single curve for BN-All-MB, since the curves with different

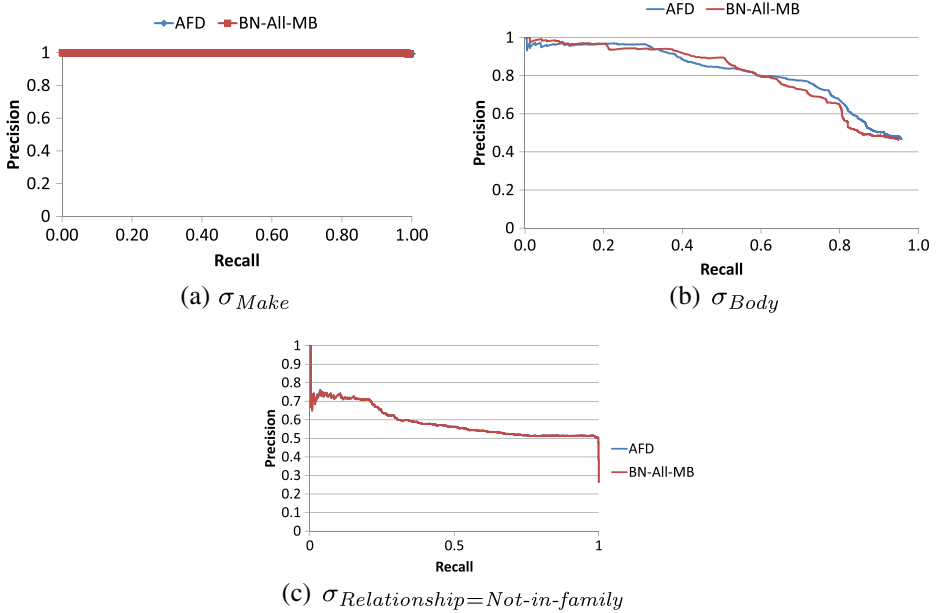


Fig. 5 Precision-recall curve for different queries

α values were coincident. This is because there are no rewritten queries with high throughput, therefore just increasing the α value does not increase recall. For BN-Beam, the level of search, L , is set to two. We see that the recall increases with increase in the value of α . Figure 6b shows the change in precision for different values of α as the number of queries sent to the database increases. As expected, the precision of BN-Beam-0.30 is higher than BN-Beam-0.35 and BN-Beam-0.40. In particular, we point out that the precision of BN-Beam-0.30 remains competitive with BN-All-MB in all the cases while providing significantly higher recall. BN-Beam

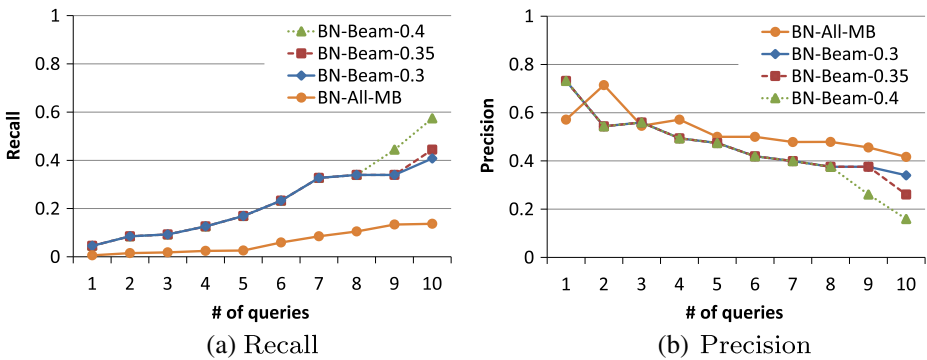


Fig. 6 Change in recall and precision for different values of α in F-measure metric for top-10 rewritten queries for $\sigma_{Year} = 2002$. The curves for all values of α for BN-All-MB are the same, so they are represented as a single curve


is able to retrieve relevant incomplete answers with high recall without any large decrease in precision: the average difference is only 0.05.

6.2.3 Comparison of multi-attribute queries


We now compare Bayesian network and AFD approaches for retrieving relevant uncertain answers with multiple missing values when multi-attribute queries are issued by the user. We note that the current QPIAD system retrieves only uncertain answers with atmost one missing value on query-constrained attributes. We compare BN-Beam with two baseline AFD approaches.

1. **AFD-All-Attributes:** This approach creates a conjunctive query by combining the best rewritten queries for each of the constrained attributes. The best rewritten queries for each attribute constrained in the original query are computed independently and new rewritten queries are generated by combining a rewritten query for each of the constrained attributes. The new queries are sent to the autonomous database in the decreasing order of the product of the expected precisions of the individual rewritten queries that were combined to form the query. AFD-All-Attributes technique is only used for multi-attribute queries where the determining set of each of the attributes are disjoint.
2. **AFD-Highest-Confidence:** This approach uses only the AFD of the query-constrained attribute with the highest confidence for generating rewritten queries, ignoring the other attributes.


We evaluate these methods for selection queries with two constrained attributes. For BN-Beam, the level of search is set to 2 and the value for α in the F-measure metric is set to zero.




(a) $\sigma_{Make=bmw \wedge Mileage=15000}$



(b) $\sigma_{Education=HS-grad \wedge Relationship=Husband}$



(c) $\sigma_{Make=kia \wedge Year=2004}$



(d) $\sigma_{WorkClass=private \wedge Relationship=own-child}$


Fig. 7 Precision-recall curve for the results returned by top-10 rewritten queries for various queries

6.2.4 Comparison of AFD-all-attributes and BN-beam


Figure 7a shows the precision-recall curve for the results returned by top ten rewritten queries by AFD-All-Attributes and BN-Beam for the query $\sigma_{Make=bmw \wedge Mileage=15000}$ issued to the car database. Figure 7b shows a similar curve for the query $\sigma_{Education=HS\text{-}grad \wedge Relationship=Husband}$ issued to the adult database. We note that the recall of the results returned by AFD-All-Attributes is significantly lower than BN-Beam in both cases (see Fig. 8). This is because the new queries generated by conjoining the rewritten queries of each constrained attribute do not capture the joint distribution of the multi-attribute query. Therefore, the throughput of these queries are often very low, in the extreme case they even generate empty queries. The precision of the results returned by AFD-All-Attributes is only slightly higher than BN-Beam (see Fig. 7). By retrieving answers with a little lesser precision and much higher recall than AFD-All-Attributes, BN-Beam technique becomes very effective in scenarios where the autonomous database has limits on the number of queries that it will respond to.

6.2.5 Comparison of AFD-highest-confidence and BN-beam


Figure 7 shows the precision-recall curves for the results returned by the top ten queries for multi-attribute queries issued to the car and adult databases. Figure 8 shows the change in recall with each of the top ten rewritten query issued to the autonomous database. We note that the recall of the results returned by AFD-Highest-Confidence is much higher than BN-Beam. However, this increase in




(a) $\sigma_{Make=bmw \wedge Mileage=15000}$



(b) $\sigma_{Education=HS\text{-}grad \wedge Relationship=Husband}$



(c) $\sigma_{Make=kia \wedge Year=2004}$



(d) $\sigma_{WorkClass=private \wedge Relationship=own\text{-}child}$

Fig. 8 Change in recall as the number of queries sent to the autonomous database increases for various queries

recall is accompanied by a drastic fall in precision. This is because AFD-Highest-Confidence approach is oblivious to the values of the other constrained attributes. Thus, this approach too, is not very effective for retrieving relevant possible answers with multiple missing values for multi-attribute queries.

7 Conclusion

We presented a comparison of cost and accuracy trade-offs of using Bayesian network models and Approximate Functional Dependencies (AFDs) for handling incompleteness in autonomous databases. We showed how a generative model of an autonomous database can be learned and used by query processors while keeping costs manageable.

We compared Bayesian networks and AFDs for imputing single and multiple missing values. We showed that Bayesian networks have a significant edge over AFDs in dealing with missing values on multiple correlated attributes and at high levels of incompleteness in test data.

Further, we presented a technique, BN-All-MB, for generating rewritten queries using Bayesian networks. We then proposed a technique, BN-Beam, to generate rewritten queries that retrieve relevant uncertain results with high precision and throughput, which becomes very important when there are limits on the number of queries that autonomous databases respond to. We showed that BN-Beam trumps AFD-based approaches for handling multi-attribute queries. BN-Beam contributes to the QPIAD system by retrieving relevant uncertain answers with multiple missing values on query-constrained attributes with high precision and recall.

References

- Batista, G., & Monard, M. (2002). A study of k-nearest neighbour as an imputation method. In *Soft computing systems: design, management and applications* (pp. 251–260). Santiago, Chile.
- Bishop, C., et al. (2006). *Pattern recognition and machine learning* (Vol. 4). Springer, New York.
- Cars.com (2013). <http://www.cars.com>. Accessed 1 Feb 2013.
- CBioC: (2013) <http://cbioc.eas.asu.edu/>. Accessed 1 Feb 2013.
- Cooper, G. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2), 393–405.
- Dempster, A., Laird, N., Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1), 1–38.
- Fernández, A., Rumí, R., Salmerón, A. (2012). Answering queries in hybrid Bayesian networks using importance sampling. *Decision Support Systems*, 53(3), 580–590.
- Frank, A., & Asuncion, A. (2010). UCI machine learning repository. <http://archive.ics.uci.edu/ml>. Accessed 1 Feb 2013.
- Geiger, D., Verma, T., Pearl, J. (1990). Identifying independence in Bayesian networks. *Networks*, 20(5), 507–534.
- Gupta, R., & Sarawagi, S. (2006). Creating probabilistic databases from information extraction models. In *VLDB* (pp. 965–976).
- Hartemink, A., et al. (2005) Banjo: Bayesian network inference with java objects. Web Site <http://www.cs.duke.edu/~amink/software/banjo/>. Accessed 1 Feb 2013.
- Heckerman, D. (1992). *The certainty-factor model* (2nd edn). Encyclopedia of Artificial Intelligence.
- Heckerman, D., Geiger, D., Chickering, D.M. (1995). Learning Bayesian networks: the combination of knowledge and statistical data. *Machine Learning*, 20(3), 197–243.
- Heitjan, D.F., & Basu, S. (1996). Distinguishing missing at random and missing completely at random. *The American Statistician*, 50(3), 207–213.

- Jensen, F., Olesen, K., Andersen, S. (2006). An algebra of Bayesian belief universes for knowledge-based systems. *Networks*, 20(5), 637–659.
- Jensen, F.V., & Nielsen, T.D. (2007). *Bayesian networks and decision graphs*. Springer.
- Khatri, H. (2006). *Query processing over incomplete autonomous web databases*. Master's thesis, Arizona State University, Tempe, USA.
- Manning, C.D., Raghavan, P., Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press.
- Minka, T., Winn, J., Guiver, J., Knowles, D. (2010). Infer.NET 2.4. <http://research.microsoft.com/infernet>. Microsoft Research Cambridge. Accessed 1 Feb 2013.
- Minka, T.P. (2001). Expectation propagation for approximate Bayesian inference. In *UAI* (pp. 362–369).
- Murphy, K., et al. (2001). The Bayes net toolbox for Matlab. *Computing Science and Statistics*, 33(2), 1024–1034.
- Muslea, I., & Lee, T. (2005). Online query relaxation via Bayesian causal structures discovery. In *Proceedings of the national conference on artificial intelligence* (Vol. 20, p. 831). Menlo Park, CA/Cambridge, MA, London: AAAI Press/MIT Press.
- Nambiar, U., & Kambhampati, S. (2006). Answering imprecise queries over autonomous web databases. In *Proceedings of the 22nd International Conference on Data Engineering (ICDE)* (pp. 45–45). IEEE.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann.
- Ramoni, M., & Sebastiani, P. (1997). Learning Bayesian networks from incomplete databases. In *Proceedings of the 13th conference on uncertainty in artificial intelligence* (pp. 401–408). Morgan Kaufmann Publishers Inc.
- Ramoni, M., & Sebastiani, P. (2001). Robust learning with missing data. *Machine Learning*, 45(2), 147–170.
- Romero, V., & Salmerón, A. (2004). Multivariate imputation of qualitative missing data using Bayesian networks. In *Soft methodology and random information systems* (pp. 605–612). Springer.
- Russell, S.J., & Norvig, P. (2010). *Artificial intelligence—a modern approach*. Pearson Education.
- Shortliffe, E. (1976) *Computer-based medical consultations: MYCIN* (Vol. 388). Elsevier, New York.
- Wolf, G., Kalavagattu, A., Khatri, H., Balakrishnan, R., Chokshi, B., Fan, J., Chen, Y., Kambhampati, S. (2009). Query processing over incomplete autonomous databases: query rewriting using learned data dependencies. *Very Large Data Bases Journal*, 18(5), 1167–1190.
- Wolf, G., Khatri, H., Chokshi, B., Fan, J., Chen, Y., Kambhampati, S. (2007). Query processing over incomplete autonomous databases. In *Proceedings of the 33rd international conference on very large data bases* (pp. 651–662). VLDB Endowment.
- Wu, C., Wun, C., Chou, H. (2004). Using association rules for completing missing data. In *Hybrid Intelligent Systems (HIS)* (pp. 236–241). IEEE.